arXiv:1610.07994v1 [math.PR] 25 Oct 2016

A note on dimers and T-graphs

Nathanaël Berestycki* Benoit Laslier† Gourab Ray‡

October 26, 2016

Abstract

The purpose of this note is to give a succinct summary of some basic properties of T-graphs which arise in the study of the dimer model. We focus in particular on the relation between the dimer model on the heaxgonal lattice with a given slope, and the behaviour of the uniform spanning tree on the associated T-graph. Together with the main result of the companion paper [1], the results here show Gaussian free field fluctuations for the height function in some dimer models.

1 Introduction

The dimer model is a classical model from statistical mechanics which can be thought of as defining a natural random surface through its height function. We refer to [2] for a detailed survey on the subject. A key tool in the early study of the dimer model is Temperley's bijection, which relates the dimer model on domains satisfying certain boundary conditions to uniform spanning trees on a modified graph. This correspondence was extended first by Kenyon, Propp and Wilson [5] to so-called Temperleyan graphs. A more delicate but also more general extension was discovered by Kenyon and Sheffield [6]. This can be used in particular to relate the dimer model on (subsets of) the hexagonal lattice to spanning trees on modified graphs called T-graphs. This correspondence is the main concern of this paper. Our goal is to make precise a number of facts which are probably known in the folklore but for which we could not find a reference, as well as to prove a number of new estimates which are needed for the study of the dimer model on the hexagonal lattice.

Our results can be summarised as follows:

1. We show that the height function of a dimer configuration is given by the winding of branches in the associated spanning tree on the T-graph (Theorem 4.4).

2. We study the large-scale behaviour of simple random walk on T-graphs. In particular, we prove recurrence (Theorem 3.10) as well as a uniform crossing estimate (Theorem 3.8).

3. We construct discrete domains U^δ on the T-graph associated with any given continuum domain D on the plane with locally connected boundary (Theorem 4.8). More precisely, given a plane $P \subset \mathbb{R}^3$ and a domain $D \subset P$ with locally connected boundary, we construct domains U^δ on an associated T-graph such that the height function of any dimer configuration on U^δ satisfies $h^\delta(\partial U^\delta) \to \partial D$ as closed sets in $\mathbb{R}^3$ when $\delta \to 0$.

4. Finally, we prove in Theorem 4.10 that the correspondence between dimers and uniform spanning trees on finite portions of the hexagonal lattice considered by Kenyon and Sheffield [6] extend to the full plane. More precisely, for a given plane $P \subset \mathbb{R}^3$ as above, the local limits of uniform spanning trees on the associated T-graph, corresponds to whole plane dimer configurations on the hexagonal lattice whose law is the unique ergodic Gibbs measure on dimers with that slope constructed by Sheffield [11].

*Supported in part by EPSRC grants EP/L018896/1 and EP/I03372X/1
†Supported in part by EPSRC grant EP/I03372X/1
‡Supported in part by EPSRC grant EP/I03372X/1

Equivalently, the weak limit of uniform spanning trees on suitable large portions of T-graphs (which exists by a result in the companion paper [1]), is in 'bijection' with the whole plane Gibbs measure on dimers of the corresponding slope.

Together with the main result of [1] (Theorem 1.2 in that paper), this implies that the height function of dimer configurations on the hexagonal lattice with planar boundary conditions (see item 3 above), has fluctuations given in the scaling by the Gaussian free field, as stated precisely in Theorem 1.1 of [1].

2 Definition and construction

2.1 Lozenge tilings and height function

We will first give a short introduction to the lozenge tiling model to setup our notations and make the paper more self-contained. For more details see [8, 2] for example.

We call **lozenge tiling** of a domain $D \subset \mathbb{C}$ a tiling of D by tiles which are made by gluing two equilateral triangles with edge length 1 (see Figure 1). Lozenge tilings are in bijection with perfect matchings of subsets of the hexagonal lattice and with surfaces in $\mathbb{R}^3$ formed by the boundary of stacks of unit cubes (step surfaces). These bijections should be visually evident from Figure 1 and are standard so we will not give more details on them. Using these bijections, we will freely identify tilings, perfect matchings and step surfaces. When talking about perfect matchings, we will write $\mathcal{H}$ for the infinite hexagonal lattice and G for subgraphs of $\mathcal{H}$. We will call one of the bipartite class of vertices in $\mathcal{H}$ black and the other white and write b and w for vertices in these classes respectively.

When D is bounded, we are interested in the uniform measure on lozenge tiling of D. When D is not bounded, one cannot simply take the uniform measure but the case $D = \mathbb{C}$ is well understood via the following theorem.

Theorem 2.1. *[11] For all p_a, p_b, p_c in $(0,1)$ such that $p_a + p_b + p_c = 1$, there exists a unique ergodic translation invariant Gibbs measure μ on lozenge tiling coverings of the plane such that*

- *Vertical (resp. north east- south west, north west-south east) lozenges appear with probability p_a (resp. p_b, p_c).*
- *For any finite subgraph $G \subset \mathcal{H}$, the measure on G, conditioned on the state of $\mathcal{H} \setminus G$, is uniform over all possible dimer configurations in G.*

The translation invariance implies directly that the expected height change is linear. Correlations between tiles in these measures are known to be expressed in terms of determinants with an explicit (and quite simple) kernel ([4]) but we will not use this fact here.

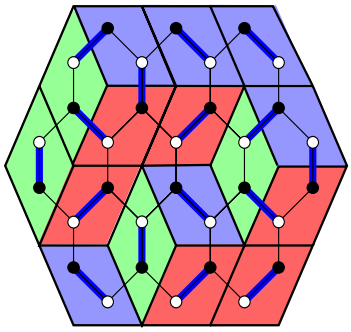

Figure 1: Lozenge tiling of a domain. The colouring of the lozenges make the bijection with stepped surfaces apparent. The height function indicated is the z coordinate of the corresponding point in $\mathbb{Z}^3$.

The height function with respect to a plane $\mathcal{P}$ in $\mathbb{R}^3$ at some vertex x of the tiling is defined to be the distance of the corresponding point of the surface to $\mathcal{P}$ (up to a global scaling factor). For example Figure 1 shows the height function with respect to the horizontal plane $z = 0$ (using the standard coordinates in $\mathbb{R}^3$).

The above definition can be made more algebraic by giving the following equivalent definition. Let ϕ_{ref} be a function on oriented edges such that $\phi_{\text{ref}}(bw) = -\phi_{\text{ref}}(wb)$ for all edges bw and such that

$$\forall w, \sum_{b \sim w} \phi_{\text{ref}}(wb) = 1 \qquad \forall b, \sum_{w \sim b} \phi_{\text{ref}}(bw) = -1.$$

To make statement about planarity simpler, we will also always assume that ϕ_{ref} is periodic even though it is not strictly required for this construction to hold. This function ϕ_{ref} is called the reference flow and in the previous geometric definition each reference plane corresponds to a particular choice of ϕ_{ref}. Given a tiling, let M be the set of edges of $\mathcal{H}$ that correspond to a lozenge (the blue edges in Figure 1). We define a flow ϕ_M on oriented edges by $\phi_M(wb) = \mathbb{1}_{(wb)\in M}$ and $\phi_M(bw) = -\mathbb{1}_{(wb)\in M}$. We now define the dual flow $(\phi_M - \phi_{\text{ref}})^\dagger$ on edges of the dual of $\mathcal{H}$ (which is a triangular lattice). The flow $(\phi_M - \phi_{\text{ref}})^\dagger$ on an oriented dual edge $e^\dagger$ crossing an oriented edge e in the primal (so that $e^\dagger$ is just a rotation of e by $\pi/2$) is the same as that in e. Since $(\phi_M - \phi_{\text{ref}})$ is divergence free by definition, we see that $(\phi_M - \phi_{\text{ref}})^\dagger$ is a gradient flow so it has a primitive h_M and we define h_M to be the height function of M with reference flow ϕ_{ref}. The global additive constant is fixed arbitrarily. A convenient choice when we work in a bounded domain is to set $h = 0$ on some fixed boundary point.

It is easy to see that the height function determines the tiling (assuming the reference is known). One can also check (it is quite clear from Figure 1) that in a bounded domain D, the height function h_M along the boundary of D is independent of M (as long as M is a tiling of D). We can therefore talk of the boundary height of D without specifying a tiling inside. We say that a domain has **planar boundary with width** C if there exists a linear function f such that $|f - h| \leq C$ on ∂D.

2.2 T-graph construction

In this section we construct the T-graphs and state some of its geometric properties that will be needed later.

We start by defining suitable coordinates on the infinite hexagonal lattice $\mathcal{H}$. We embed the hexagonal lattice $\mathcal{H}$ in the plane in the way represented in Figure 2. The figure which is drawn with thicker lines in Figure 2 is called the fundamental domain and we write it as $\mathcal{H}_1$. We let e_1 and e_2 be the two vectors represented in Figure 2. Given v a vertex of $\mathcal{H}$, we call coordinates of v the unique $(m, n) \in \mathbb{Z}^2$ such that $v - me_1 - ne_2 \in \mathcal{H}_1$. Note that given (m, n) there are exactly two vertices with coordinates (m, n), the top one is a white vertex by convention and hence the bottom one is black. We will write $m(v)$ and $n(v)$ for the coordinates of the vertex v. We will also write $b(m, n)$ and $w(m, n)$ for the black and white vertices with coordinates (m, n).

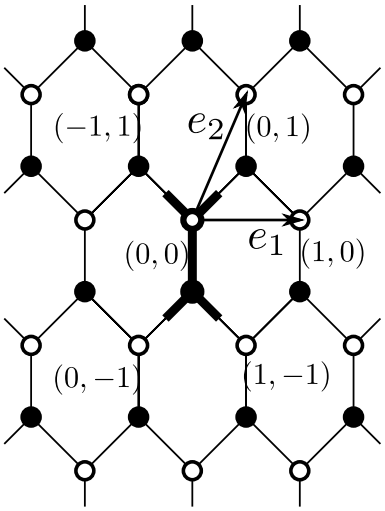

Figure 2: An illustration of the coordinates we use on the hexagonal lattice. Near each vertical edge are indicated the (common) coordinates of its two endpoints.

We write $\mathcal{H}^\dagger$ for the dual graph of $\mathcal{H}$. This is a triangular lattice. Each of its faces contains a vertex of $\mathcal{H}$ and a face is called black or white according to the colour of that vertex. Vertices of $\mathcal{H}^\dagger$ can be associated to the point in the centre of a face of $\mathcal{H}$. For a vertex v of $\mathcal{H}^\dagger$ we let $(m(v), n(v))$ be the (common) coordinates of the two vertices of the face corresponding to v located just to the right of v.

T-graphs will be defined as the primitive of a specific gradient flow on $\mathcal{H}^\dagger$ which we define now. Let us fix p_a, p_b, p_c in $(0, 1)$ such that $p_a + p_b + p_c = 1$. Let Δ be a triangle in the complex plane with angles $\pi p_a, \pi p_b, \pi p_c$. We write A, B, C its vertices and $\alpha = \vec{BC}$, $\beta = \vec{CA}$, $\gamma = \vec{AB}$ seen as complex numbers. We have $\alpha + \beta + \gamma = 0$. Let λ be a complex number of modulus one. We will require later λ to be outside a set of Lebesgue measure 0 and all references to almost every λ are with respect to the Lebesgue measure on the circle.

We can now define the following flow on oriented edges between a white vertex w and a black vertex b:

$$\phi(wb) = \Re\left(\lambda^{-1}(\tfrac{\beta}{\gamma})^{-m(w)}(\tfrac{\beta}{\alpha})^{-n(w)}\right)\alpha\lambda(\tfrac{\beta}{\gamma})^{m(b)}(\tfrac{\beta}{\alpha})^{n(b)},$$

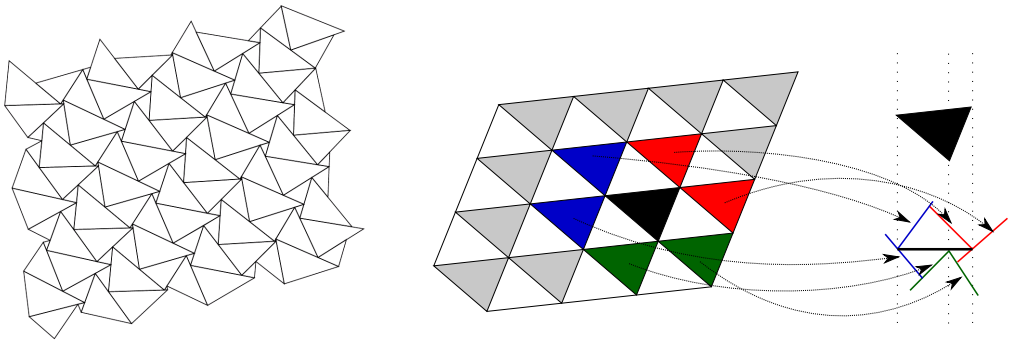

Figure 3: Left: A part of a T-graph. Right: A schematic view of the image of a black face and its neighbours in $\mathcal{H}^\dagger$ (for clarity black faces are colored differently). Each black face is projected into a segment. The vertex of the triangle which gets projected in the interior of the segment becomes a vertex of the T-graph where two other segments meet.

and $\phi(bw) = -\phi(wb)$. We then define the dual flow $\phi^\dagger$ by rotating ϕ anticlockwise by $\frac{\pi}{2}$, i.e crossing the edge wb with the white vertex on the left gives flow $+\phi(wb)$.

One can check from the formula that the circulation of the dual flow $\phi^\dagger$ around any triangle (sum of the flow along the edges in a triangular face) is 0 so $\phi^\dagger$ is a gradient flow. Hence there exists a function ψ on the vertices of the triangular lattice $\mathcal{H}^\dagger$ (we write $\psi_{\lambda\Delta}$ if we want to emphasize the dependence), unique up to an additive constant such that for all adjacent vertices of the triangular lattice $v, v' \in \mathcal{H}^\dagger$, $\psi(v') - \psi(v) = \phi^\dagger(vv')$. We fix the additive constant by specifying ψ to be 0 on the dual vertex just to the left of the fundamental domain $\mathcal{H}_1$. We can extend ψ affinely to the edges of $\mathcal{H}^\dagger$ so that ψ maps $\mathcal{H}^\dagger$ to a connected union of segments in $\mathbb{C}$ and we write $T_{\lambda\Delta} = \psi_{\lambda\Delta}(\mathcal{H}^\dagger)$ (see Figure 3 for an example). We call the image of $\mathcal{H}^\dagger$ under ψ the **T-graph** with parameters Δ and λ.

Remark 2.2. One can possibly get some intuition about the construction by comparing with what we get if we omit the real part in the definition of ϕ above. In that case all faces would have been mapped to similar triangles and we would actually have defined simply the embedding of the triangular lattice with triangles of angles πp_a, πp_b, πp_c. The T-graph can be seen as a perturbation of that embedding where triangles corresponding to black faces are flattened into segments but that otherwise preserves the adjacency structure. We will also see in Proposition 3.1 point (i) that at large scale the two embedding are close[1].

3 Properties of T-graphs

3.1 Geometric properties

The following proposition tells us how the value of λ behave under translation.

Proposition 3.1. *Let $T_{\lambda\Delta}$ be the graph constructed above (recall that we choose $\psi = 0$ on the dual vertex just to the left of the fundamental domain $\mathcal{H}_1$). Let v be the vertex of $\mathcal{H}^\dagger$ with coordinates (m, n) and let T' be the graph constructed in the same way but taking $\psi(v) = 0$. Then we have $T' = T_{\lambda'\Delta}$ with $\lambda' = \lambda(\frac{\beta}{\gamma})^m(\frac{\beta}{\alpha})^n$. In particular, if we pick λ to be uniform in the unit circle, then $T_{\lambda\Delta}$ becomes translation invariant.*

Here are some other geometric facts about T, see for instance [7] :

Proposition 3.2. *T has the following properties (see Figure 3):*

(i) Let $\ell(m, n) = \frac{\alpha}{2}m - \frac{\gamma}{2}n$, then $\psi(v) - \ell(m(v), n(v))$ is bounded. Furthermore ℓ is invertible so any point $z \in \mathbb{C}$ is at bounded distance from T.

[1] up to global factor $\frac{1}{2}$

(ii) *The image of any black face of $\mathcal{H}^\dagger$ is a segment.*

(iii) *The image of any white face is a contraction and rotation of Δ. In particular, the map preserves orientation.*

(iv) *The length of segments are bounded above and below uniformly in λ and for almost every λ no triangle is degenerate to a point.*

(v) *For any λ (resp. almost every λ), for any vertex v of $\mathcal{H}^\dagger$, $\psi(v)$ belongs to at least (resp. exactly) three segments. Generically any vertex is an endpoint of two segments and is in the interior of the third one. All endpoints of segments are of the above form $\psi(v)$ with v a vertex of $\mathcal{H}^\dagger$. We call these points* ***vertices of the T-graph****.*

(vi) *The triangular images of white faces cover the plane and do not intersect, that is, any $z \in \mathbb{C}$ not in a segment belongs to a unique face of the T-graph.*

(vii) *If two segments intersect, the intersection point is an endpoint of at least one of the segments.*

From now on we assume that λ is chosen so that no triangle is degenerate and all vertices belong to exactly three segments. We say that the corresponding T-graph is **non-degenerate**.

We can also naturally define **edges of the T-graph** to be the portion of the segments joining two vertices and faces to be the connected components of the complement of the segments. From the Proposition 3.2, we see the following correspondences.

T-graph	Hexagonal lattice	Triangular lattice
Segment	Black vertex	Black face
Face	White vertex	White face
Vertex	Face	Vertex
Edge		Edge

Note that the first three correspondences are really bijective while in the last case, some edges of $\mathcal{H}^\dagger$ do not correspond to any edge of T. With an abuse of notation, we will allow ourself to write $\psi(w)$ and $\psi(b)$ for the corresponding triangle or segment. We will also write ψ^{-1} for the inverse of this one to one correspondence.

Thanks to the above abuse of notation, ψ describes an embedding of $\mathcal{H}$ in the plane:

Corollary 3.3. *The following embedding of $\mathcal{H}$ in the plane is proper, i.e. has non crossing edges: each white vertex w is placed in the centre of mass of the triangle $\psi(w)$, each black vertex b is placed at the unique vertex of T in the interior of $\psi(b)$. Draw edges between a white vertex w and a black vertex b if the segment corresponding to b share more than one point with the face corresponding to w in the T-graph.*

Proof. Clearly a midpoint of a segment can be identified with a black face of the triangular lattice. We now simply observe that around each such vertex we drawn a hexagon:

It is clear that we described an embedding of $\mathcal{H}$ so we only have to check that edges are non crossing. For this, note that any edge (bw) lies inside the triangle $\psi(w)$. Since triangles do not intersect, no pair of edges (bw), $(b'w')$ with $w \neq w'$ intersects. Now for any w, its three adjacent edges are straight lines towards different points of $\psi(w)$ so they do not intersect. Further, the segment corresponding to a black vertex which is not a neighbour of the white vertex w in the hexagonal lattice share at most one point with the face of w. The proof is now complete. □

Finally one can associate sub-domains of a T-graph to sub domains of $\mathcal{H}$ even though it require some care on the boundary.

Proposition 3.4. *Let C be a simple (unoriented) loop on a T-graph T, then $\psi^{-1}(C)$ is a simple close curve on $\mathcal{H}^\dagger$ (each interior point actually has two pre-images, we use the one lying on the edge $e \in \mathcal{H}^\dagger$ where e is the edge which is not mapped to the whole segment). Let U be the subgraph of $\mathcal{H}$ inside $\psi^{-1}(C)$ and let D be the open domain strictly inside C. We have*

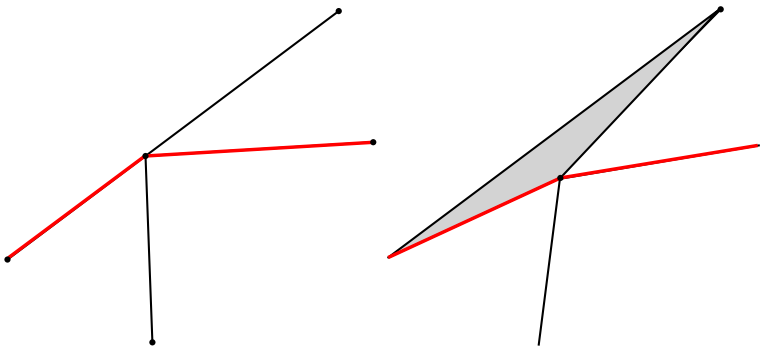

Figure 4: Thickening of an edge into a black triangle, and effect on loop.

- *The white vertices of U are exactly the w such that $\psi(w) \subset \bar{D}$.*
- *Any segment of T which is strictly inside C except possibly at finitely many points corresponds to a black vertex of U. All the neighbours of such a vertex are in U.*
- *For a segment S of T with an infinite intersection with C, consider the side of S where it has a single adjacent triangle. The black vertex $\psi^{-1}(S)$ is in U if and only if that side of S is inside C. These vertices have at least one neighbour in U and one outside U.*

Proof. Firstly, it is easy to check that $\psi^{-1}(C)$ is a simple loop on $\mathcal{H}^\dagger$. Indeed with our convention, ψ^{-1} is injective and edges including their endpoints in the T-graph are sent to edges including their endpoints in the triangular lattice.

We can clearly define a perturbation of ψ such that black triangles are sent to thin triangles rather than line segments and such that it respects the structure of the T-graph otherwise. This perturbation defines a proper embedding of $\mathcal{H}^\dagger$ and $\mathcal{H}$, therefore we can identify U by looking at the interior of $\psi^{-1}(C)$ in that embedding. In that embedding, the first two points are immediate. The last one is clear from the picture of the perturbed embedding in the bottom-right panel of Figure 5. □

3.2 Uniform crossing estimate

Definition 3.5. *Let T be a non-degenerate T-graph. The random walk X_t on T is the continuous time Markov process on the vertices of T defined by the following jump rates. If the process is at a vertex v of T, call v^+, v^- the endpoints of the unique segment which contains v in its interior. The rates of the jumps from v to $v^\pm$ are $1/|v^\pm - v|$.*

Note that this random walk is automatically a martingale thanks to the choice of the jump rates. The jump rates defined above allow us to consider a T-graph as a weighted oriented graph. From now on we view a T-graph either as a weighted directed graph or as a subset of $\mathbb{C}$ as required by context. We denote by X^v the walk started in v.

We now prove the Russo-Seymour-Welsh type uniform crossing estimate. The key input is the following uniform ellipticity bound.

Lemma 3.6 ([7] Proposition 2.22)**.** *There exists $c_0 > 0$ (depending continuously on Δ) such that for any $\mathbf{n} \in \mathbb{S}^1$ and any $v \in T$,*

$$c_0 < \mathrm{Var}(X_1^v \cdot \mathbf{n}) < c_0^{-1}.$$

Using this bound we first control the angle at which the random walk exits a ball. Let T^δ denote the T-graph rescaled by δ. We emphasise that the random walk keeps the original unrescaled time parametrisation.

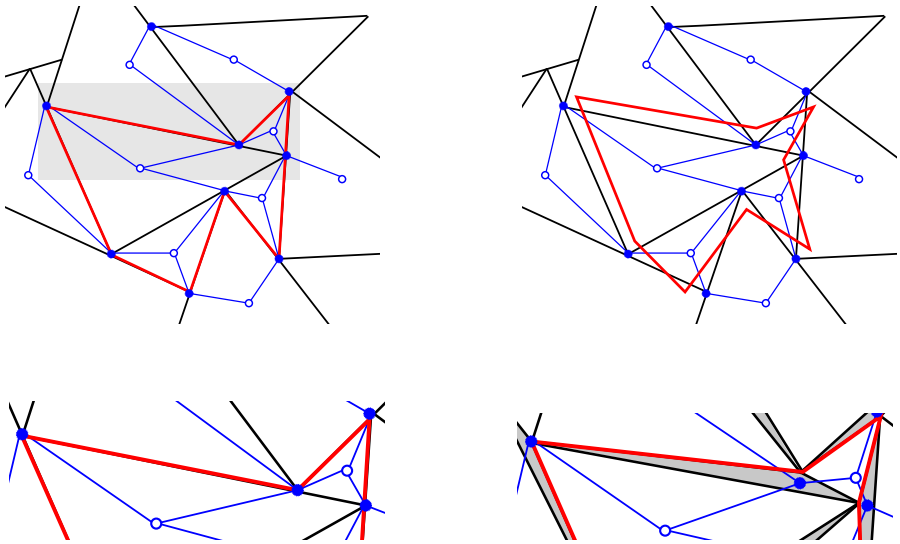

Figure 5: Top-Left: A portion of a T-graph with a loop C in red along with the embedding of the hexagonal lattice determined by the T-graph as in Corollary 3.3. Top-Right: To visualise the path $\psi^{-1}(C)$, one can shift every vertex of C in the direction where 3 triangles meet at the vertex (into the middle triangle, to be more precise). This path is drawn and can be readily seen as a dual path of the hexagonal lattice. Bottom-Left : The greyed area from the top-left figure zoomed in. Bottom-Right : Another way to visualise $\psi^{-1}(C)$: a perturbation of the T-graph where black faces are sent to thin triangles. The red path on that perturbation clearly coincide with the one in the top-left part.

Lemma 3.7. *There exists δ_1 and η such that for all $\delta \leq \delta_1$, for any $v \in T^\delta$, if τ denotes the exit time of $B(v,1)$ then for all $\theta \in [-\pi, \pi)$,*

$$\mathbb{P}\left[\mathrm{Arg}(X_\tau^v - v) \in [\theta, \theta + \pi - \eta]\right] > \eta$$

where Arg *denote the usual argument in $[-\pi, \pi)$ but the interval $[\theta, \theta + \pi - \eta]$ is interpreted cyclically if $\theta + \pi - \eta \geq \pi$.*

Proof. First let us prove that we can choose η such that for all $\theta \in [-\pi, \pi)$

$$\mathbb{P}\left[\arg(X_\tau^v - v) \in [\theta, \theta + \pi - \eta]\right] > 0. \tag{3.1}$$

Let us translate and rotate the coordinate frame to set $v = 0$ and $\theta = \eta/2$ so that the arc $[\theta, \theta + \pi - \eta]$ is symmetric around the vertical axis. We let Y_t denote the projection of the random walk on the vertical axis.

Since the random walk is a martingale we have $\mathbb{E}[Y_\tau] = 0$ by optional stopping (which we can use since $Y_{\tau \wedge t}$ is uniformly bounded by 1). Further Lemma 3.6 ensures that $\mathrm{Var}(Y_\tau) \geq cc_0$ where c_0 is as in Lemma 3.6. To see this first notice from Lemma 3.6 that $(Y_k^2 - c_0\delta^2 k)_{k \in \mathbb{N}}$ is a submartingale. Further notice that by Burkholder–Davis–Gundy inequality with $p = 2$ we have $\mathbb{E}(\tau) \geq c\delta^{-2}$. Using this, the optional stopping theorem and the monotone and dominated convergence theorems,

$$\mathbb{E}(Y_\tau^2) \geq c_0\delta^2\mathbb{E}(\tau) \geq cc_0. \tag{3.2}$$

Suppose by contradiction that

$$\mathbb{P}\left[\arg(X_\tau^v - v) \in [\theta, \theta + \pi - \eta]\right] = 0.$$

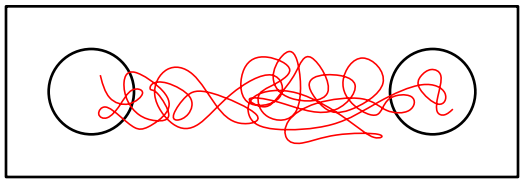

Figure 6: An illustration of the uniform crossing property.

Then Y_τ is a random variable in $[-1+O(\delta), \sin(\eta/2)+O(\delta)]$. Since $\mathbb{E}[Y_\tau]=0$, we have a bound $\mathrm{Var}(Y_\tau) \leq (\sin(\eta/2)+O(\delta))(1+O(\delta))$. Indeed, this bound is valid for any centred random variable in $[-1-O(\delta), \sin(\eta/2)+O(\delta)]$. However we saw above that $\mathrm{Var}(Y_\tau) \geq cc_0$. This is a contradiction if η, δ are small enough.

To obtain a lower bound of at least η in (3.1), just observe that there is also a universal bound of the variance of variables in $[-1,1]$ with $\mathbb{P}[Y > \sin(\eta/2)] \leq \eta$ and $\mathbb{E}[Y]=0$ so the same proof applies. □

We now give a precise definition of what it means for a graph G (embedded in $\mathbb{C}$) to satisfy the **uniform crossing property**. Let G^δ denote the rescaling of the edges of the graph G by δ. Let A^δ denote the graph spanned by the vertices of G^δ in A for any $A \subset \mathbb{C}$. Let R be the horizontal rectangle $[0,3]\times[0,1]$ and R' be the vertical rectangle $[0,1]\times[0,3]$. Let $B_1 := B((1/2,1/2),1/4)$ be the *starting ball* and $B_2 := B((5/2,1/2),1/4)$ be the *target ball* (see Figure 6). There exists universal constants $\delta_0 > 0$ and $\alpha_0 > 0$ such that for all $z \in \mathbb{C}$, $\delta \in (0,\delta_0)$, $v \in B_1$ such that $v+z \in G^\delta$,

$$\mathbb{P}_{v+z}(X \text{ hits } B_2+z \text{ before exiting } (R+z)^\delta) > \alpha_0. \tag{3.3}$$

The same statement as above holds for crossing from right to left, i.e., for any $v \in B_2$, (3.3) holds if we replace B_2 by B_1. Also, the same statement holds for the vertical rectangle R'. Let $B_2' = B((1/2,5/2),1/4)$. Then for all $z \in \mathbb{C}$, $\delta \in (0,\delta_0)$, $v \in B_1$ such that $v+z \in G^{\#\delta}$,

$$\mathbb{P}_{v+z}(X \text{ hits } B_2'+z \text{ before exiting } (R'+z)^\delta) > \alpha_0.$$

Again, the same statement holds for crossing from top to bottom, i.e., from B_2' to B_1.

Theorem 3.8. *Any non degenerate T-graph satisfies the uniform crossing property.*

Proof. We will only prove (3.3) as the proof of the other cases are identical. The idea is to look at the random walk stopped when it exits small macroscopic discs. We then construct a function on the rectangle which has a positive probability to decrease at each of these steps and such that if it decreases for the k first steps (where k is uniform in the location of the rectangle) a crossing has to happen.

Let $f(x,y) = -x + 10(y-\frac{1}{2})^2 + [10(x-\frac{5}{2}) \vee 0]^{10}$. Notice that f achieves its global minimum in B_2 and its value on ∂R is bigger than in B_1. Furthermore, it has bounded second derivatives in $R \setminus B_2$ and the norm of its gradient is also lower bounded in $R \setminus B_2$. Choose η as in Lemma 3.7. There exist $r, \eta' > 0$ such that

$$\forall z \in R \setminus B_2,\ \exists \theta \ s.t\ \forall\ \varphi \in [\theta, \theta+\pi-\eta],\ f(z+re^{i\varphi}) \leq f(z) - \eta'.$$

This is easily seen from writing a Taylor expansion around z, choosing $\theta + \pi/2 - \eta/2$ to be the direction opposite to the gradient and choosing r, η' small enough. We fix such a pair r, η'. Then we define a sequence of stopping time τ_k by $\tau_0 = 0$ and $\tau_{k+1} = \inf\{t > \tau_k | X_t \notin B(X_{\tau_k}, r)\}$. Lemma 3.7 and the definition of r, η' above show that until the random walk exits $R \setminus B_2$, it satisfies $\mathbb{P}[f(X_{\tau_{k+1}}) \leq f(X_{\tau_k}) - \eta'] \geq \eta$. Call a step good if $f(X_{\tau_{k+1}}) \leq f(X_{\tau_k}) - \eta'$. Since the value of f on ∂R is bigger than that in B_1, if the first $O(1/\eta')$ steps are all good then the random walk must hit B_2 before exiting R. Clearly this has positive probability so we are done. □

3.3 Recurrence

In this section, we prove that the random walk on a T-graph is recurrent. The main point is an asymptotic estimate of the conjugate Green function (i.e the discrete version of argument function). Let us note that the proof of this estimate comes from the exact solvability of the dimer model.

Proposition 3.9. *[3] Let T be a T-graph of parameters λ, Δ. Let w be a face of T and let d be a half line from the interior of w to infinity that avoids all vertices of T. There exists a unique (up to a constant) function $G^*_{wd} : T \to \mathbb{C}$ that is discrete harmonic except for a discontinuity when crossing d: that is, if q_{xy} are the rates of the continuous time simple random walk on T then*

$$\sum_{y\sim x} q_{xy}(f(y)-f(x)) = \#\{y\sim x : (x,y) \text{ crosses } d \text{ clockwise }\} - \#\{y\sim x : (x,y) \text{ crosses } d \text{ anticlockwise }\}$$

This function satisfies

$$G^*_{wd}(\psi(b)) = \frac{1}{2\pi}\Big(\arg_d(\psi(b)-w) + \frac{\Im\left(\lambda(\frac{\beta}{\gamma})^{m(w)}(\frac{\beta}{\alpha})^{n(w)}\right)}{\Re\left(\lambda(\frac{\beta}{\gamma})^{m(w)}(\frac{\beta}{\alpha})^{n(w)}\right)}\log|\psi(b)-w|\Big) + C + O(1/|\psi(b)-w|)$$

where $\arg_d$ denotes the determination of the argument with a 2π discontinuity on the half line d, α,β,γ are the angles of Δ and C is the constant up to which G^*_{wd} is defined.

In fact, note that because T is a T-graph, at most one neighbour of x is such that (x,y) crosses d. Note that by definition, if X_t is a random walk on T, then $G^*_{wd}(X_t) - \#(\text{signed crossings of } d \text{ before time } t)$ is a martingale. From this it is easy to check recurrence.

Theorem 3.10. *The random walk on a T-graph is recurrent.*

Proof. Let T be a T-graph and let e be an (oriented) edge of T to be chosen appropriately later. Let w^+ and w^- denote the triangles on the left and right of e. Let X_t denote the random walk started at the end of e, let V_t denote the number of jumps along e before time t. For $R>0$, let τ_R denote the exit time from the ball of radius R centered at e. Let d^+ be a half line starting in w^+, crossing e to go into w^- and avoiding all vertices of T. Let d^- be a piece of d^+ starting in w^- and going to infinity. for ease of notation, we write $G^+ = G^*_{w^+d^+}$, $G^- = G^*_{w^-d^-}$.

Note that we can always chose e so that

$$\frac{\Im\left(\lambda(\frac{\beta}{\gamma})^{m(w^+)}(\frac{\beta}{\alpha})^{n(w^+)}\right)}{\Re\left(\lambda(\frac{\beta}{\gamma})^{m(w^+)}(\frac{\beta}{\alpha})^{n(w^+)}\right)} \neq \frac{\Im\left(\lambda(\frac{\beta}{\gamma})^{m(w^-)}(\frac{\beta}{\alpha})^{n(w^-)}\right)}{\Re\left(\lambda(\frac{\beta}{\gamma})^{m(w^-)}(\frac{\beta}{\alpha})^{n(w^-)}\right)}.$$

We note that $V_t - G^+(X_t) + G^-(X_t)$ is a martingale. Further, we have $G^+(X_{\tau_R}) - G^-(X_{\tau_R}) = C\log R + O(1)$ for some C. Therefore $\mathbb{E}V_{\tau_R} = C\log R + O(1)$ which implies that the probability of reaching distance R before returning to e decays like $O(1/\log R)$. This implies recurrence of the endpoint of e. Since a nondegenerate T-graph is irreducible (Lemma 3.23 in [7]), the whole graph is recurrent. □

4 Dimers and UST

4.1 Definition of the mapping

In this section we describe the mapping from a forest on T to a perfect matching of $\mathcal{H}$ (a subset of edges of $\mathcal{H}$ such that every vertex is incident to exactly one edge, i.e., a dimer configuration). This mapping was defined in [6] actually in a more general setting. We first give the construction in a full plane setting before addressing finite domains.

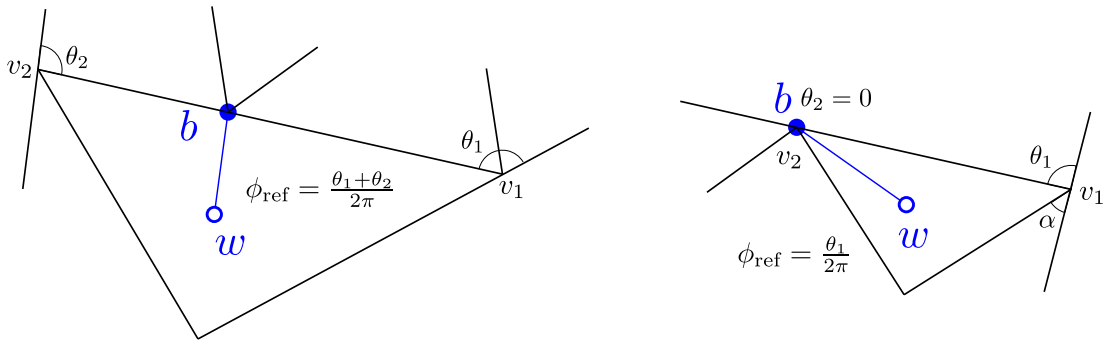

Figure 7: The reference flow associated to a T-graph. The left hand side shows a case where S_b is different from S_1 and S_2 while the right hand side shows a case with $S_2 = S_b$. Note that the local structures at v_1 and v_2 can be of two possible types as seen in the left hand side. In particular the angle α in the right hand side could be equal to 0.

Let T be a fixed T-graph and let F be a spanning forest on T. In the context of an oriented graph as here, this means the following: every vertex has a single outgoing edge in F and there is no loop in F, even ignoring the orientation. The forest F is not necessarily connected and each connected component of F is called a tree. Trees in F naturally inherit the orientation from the T-graph T and it is easy to see that each tree has a unique branch oriented toward infinity. Let $F^\dagger$ denote the dual of F, $F^\dagger$ is also a spanning forest with no finite connected component. For each connected component of $F^\dagger$ we choose one end and orient edges towards that end. Observe that if F is a one-ended tree then $F^\dagger$ is also a one ended tree and there is therefore no arbitrary choice in the orientation of $F^\dagger$.

Definition 4.1. *Given $F^\dagger$ and its orientation we can construct a dimer configuration as follows. To each white vertex w of $\mathcal{H}$, is associated a face $\psi(w)$ of T. In $F^\dagger$, there is a unique outgoing edge starting in $\psi(w)$. This edge crosses an edge part of a segment of T and by construction this segment corresponds to some black vertex $b \in \mathcal{H}$ which is adjacent to w in $\mathcal{H}$ (see Corollary 3.3). We define this vertex b to be the match of w. We denote by $M(F^\dagger)$ the matching constructed by the above procedure. By abuse of notation when F is a one ended tree we write $M(F)$ for $M(F^\dagger)$ with its unique orientation.*

Proposition 4.2. *Let F be a spanning forest of a non-degenerate T-graph and let $F^\dagger$ denote an orientation of its dual. In the above construction, $M(F^\dagger)$ is a perfect matching of $\mathcal{H}$.*

Proof. It is clear that every white vertex is matched by construction to a single black vertex so we only need to check that no black vertex is matched with two white vertices. Let b be a black vertex and let $v \in T$ be the unique vertex in the interior of the segment $\psi(b)$ associated to b. There is exactly one outgoing edge from v in F so $\psi(b)$ is crossed by exactly one edge in $F^\dagger$ which of course has only one orientation. By construction this implies that b is matched to only one white vertex. □

The main quantity of interest in a dimer configuration is often the height function, which completely encodes the dimer configuration. We will also explain below how it can be related to the winding of branches in the associated forest F. However, recall from Section 2.1 that the definition of the height involves an arbitrary choice of a "reference flow"; we therefore first describe how to choose this flow in order to have an exact identity.

The right choice of reference flow was introduced in [3] and is defined geometrically as in Figure 7. Let w and b be adjacent white and black vertices. Let v_1 and v_2 be the two vertices of T on both sides of $\psi(w) \cap \psi(b)$ and let S_1 and S_2 be the two segments containing v_1 and v_2 in their interior respectively (note that one of them might be $\psi(b)$). The flow from w to b is defined to be $\phi_{\text{ref}}(wb) = \frac{\theta_1+\theta_2}{2\pi}$ where θ_1 (resp. θ_2) is the angle between $\psi(b)$ and S_1 (resp. S_2) measured opposite to $\psi(w)$ and without sign.

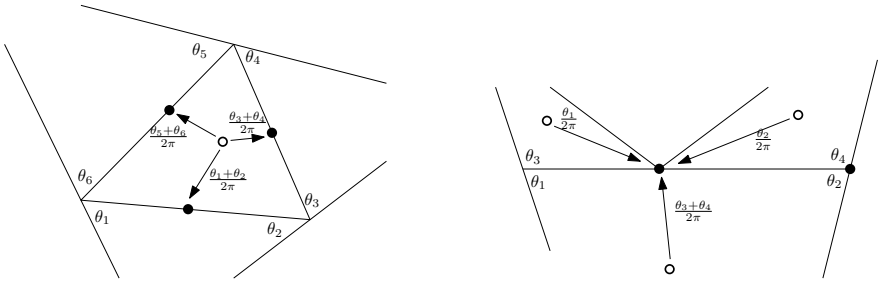

Figure 8: An illustration of why ϕ_{ref} is a valid reference flow. In the left figure, some of the angles might be 0.

Proposition 4.3. *The reference flow satisfies for every white vertex $w \in \mathcal{H}$, $\sum_{b\sim w} \phi_{ref}(wb) = 1$, and for every black vertex $b \in \mathcal{H}$, $\sum_{w\sim b} \phi_{ref}(bw) = -1$.*

The proof of the proposition follows from Figure 8 which illustrates the generic situation for a face in a T-graph, note that some of the angles however may be equal to zero. Proposition 4.3 says that the divergence of ϕ_{ref} is ± 1 at every white/black vertex; as we will see in a moment, this property allows us to define a function for a dimer configuration unambiguously and hence we speak of a *reference flow*.

Let M be a dimer configuration on $\mathcal{H}$ (seen as a subset of edges) and recall the definition of the height function: We define a flow ϕ_M by $\phi_M(wb) = -\phi_M(bw) = 1$ if $(wb) \in M$ and 0 otherwise. Then ϕ_M is also a flow and satisfies the same property as ϕ_{ref} in Proposition 4.3. Hence $\phi_M - \phi_{\text{ref}}$ is a divergence free flow. Consequently there exists a unique (up to constant) function h defined on the vertices of $\mathcal{H}^\dagger$ such that if $e = bw$ is an edge of $\mathcal{H}$ and $e^\dagger = xy$ is the corresponding dual oriented edge in $\mathcal{H}^\dagger$ then

$$(\phi_M - \phi_{\text{ref}})(wb) = h(y) - h(x).$$

h is called the **height function** of the (infinite) dimer configuration M.

The main result of this section relates the height differences to winding of branches when the dimer configuration derives from a one-ended spanning tree. Let F be a spanning tree of T. Let v, v' be two vertices of $\mathcal{H}^\dagger$. Let $\gamma_{v\to v'}$ be the unique path in F joining v and v'. That is, $\gamma_{v\to v'} = (\gamma_0, \gamma_1, \dots, \gamma_n)$ where $\gamma_0 = v$ and $\gamma_n = v'$ and the unoriented edge $\{\gamma_i, \gamma_{i+1}\} \in T$. To this we add two vertices γ_{-1} and γ_{n+1} as follows: let S_0 (resp. S_n) be the segment of which γ_0 (resp γ_n) is the midpoint. We require (γ_{-1}, γ_0) to be perpendicular to S_0 and γ_{-1} is on the side of S_0 containing the two incoming edges to γ_0. Likewise, we require (γ_n, γ_{n+1}) to be perpendicular to S_n and γ_{n+1} is on the side of S_n not containing the incoming edges to γ_n. See Figure 9 for an illustration.

Theorem 4.4. *Let T be a non degenerate T-graph and let F be a one ended spanning tree of T. Let h denote the height function of $M(F)$ with the reference flow ϕ_{ref} defined above. We have for any $v, v' \in \mathcal{H}^\dagger$,*

$$h(v') - h(v) = \frac{1}{2\pi} W_{\text{int}}(\gamma_{v\to v'})$$

where $\gamma_{v\to v'}$ is the path in F joining v and v' as defined above, and W_{int} denotes the sum of the angles turned (with signs) from the initial point γ_{-1} to final point γ_{n+1} of the path.

Proof. First note that with this definition $W_{\text{int}}(\gamma_{v\to v'})$ is additive and antisymmetric (in both cases this is not completely obvious because of the beginning and end portions added to the path); see Figure 9. By additivity, we can restrict ourself to the case where $\psi(v)$ and $\psi(v')$ are joined by a single edge in the tree. Furthermore by antisymmetry we can assume without loss of generality that the edge in the tree is oriented from $\psi(v)$ to $\psi(v')$ as in the left hand side of Figure 9.

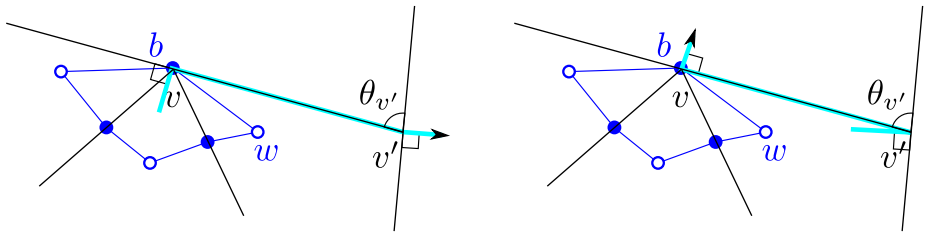

Figure 9: Illustration of the proof of Theorem 4.4. The hexagonal lattice is shown in blue. Left: The cyan path is $\gamma_{v\to v'}$ and in this case its intrinsic winding is clearly $-\frac{\pi}{2} - \theta_{v'} + \frac{\pi}{2} = -\theta_{v'}$. Right: The cyan path is $\gamma_{v'\to v}$ and the winding is $\theta_{v'}$.

Let b be the black vertex associated to v. Since there is an edge from $\psi(v)$ to $\psi(v')$ (which we have assumed is oriented from $\psi(v)$ to $\psi(v')$), it must be that $\psi(v')$ is an endpoint of $\psi(b)$. So v and v' are neighbours in $\mathcal{H}^\dagger$ (this is obvious if we think of the segment as a flattened triangle, where the edge between the midpoint and the extremity is just one of the edges of the triangle). Let w be the white vertex of the edge e of $\mathcal{H}$ which is dual to the edge between v and v'. Note that the black vertex of e is b. From the construction of the T-graph, we see easily that $\psi(w) \cap \psi(b)$ is the edge $[\psi(v), \psi(v')]$ of T. Since that edge is in F, the outgoing edge from $\psi(w)$ in $F^\dagger$ does not cross $\psi(b)$. Hence b and w are not matched together and

$$h(v') - h(v) = \pm\phi_{\text{ref}}(wb)$$

Recall that $|\phi_{\text{ref}}(wb)| = \frac{\theta_v + \theta_{v'}}{2\pi}$ where θ_v and $\theta_{v'}$ are angles at v and v'. However by definition b is the black vertex associated to v so $\theta_v = 0$. Also, it is immediate from Figure 9 that $|W_{\text{int}}(\gamma_{v\to v'})| = \theta_{v'}$. It remains only to check that the signs are consistent with our conventions. For this, note that since the orientation of black triangles are preserved, if the tree edge $[\psi(v), \psi(v')]$ goes to the right (resp. left) then the edge vv' is crossed with the white vertex to its right (resp. left). Therefore the sign of $\phi_{\text{ref}}(vv')$ is determined by this turn and it is easy to check that it is opposite to the sign of $W_{\text{int}}(\gamma_{v\to v'})$. □

Remark 4.5. Note that we add two auxilary segments at the initial and final points of the paths (the segments (γ_{-1}, γ_0) and (γ_n, γ_{n+1})) which has a contribution to the winding which we did not include in the definition of Theorem 5.1 in [1]. However, we can add this to term $m^{\#\delta}$ there since this extra winding is strictly a function of the point in the T-graph and not of the uniform spanning tree.

Remark 4.6. There exists a similar construction of T-graphs for dimers on the square lattice and actually even any periodic bipartite planar graph and a version of Theorem 4.4 holds in that more general setting (see [6]). We do not work in this more general setup because the Central Limit Theorem of [7], which is needed for the convergence of the height function to the Gaussian free field in [1] is currently only proved for the hexagonal lattice (although the proof would extend to the square lattice without major difficulties).

Remark 4.7. The reference flow ϕ_{ref} is essentially equivalent to the flow $\phi_{p_a p_b p_c}$ defined as $\phi_{p_a p_b p_c}(wb) = p_a$ (resp. p_b, p_c) if wb is a vertical (resp. NE-SW, NW-SE) edge. Indeed it follows from [3] Section 3.2.3 that if h and h' are the two heights functions associated to the same tiling but with references ϕ_{ref} or $\phi_{p_a p_b p_c}$ then $|h - h'| \leq C$ for some universal constant C.

4.2 Dimer configurations on finite domains

In the previous section we explained how to associate to a spanning tree of an (infinite) T-graph a dimer configuration on the whole hexagonal lattice. We now explain how to extend this construction to spanning trees of finite subgraphs of a T-graph for a given finite domain of the hexagonal lattice. This will require choosing the boundary of the discrete domain on the T-graph in a careful way. In this section, we assume

that we are given a continuous domain and that we want to approximate it by a"good" subgraph of the hexagonal lattice.

Let $U \subset \mathbb{C}$ be a bounded domain such that ∂U is locally connected. Let $x \in \partial U$ be a marked point on the boundary. Consider an (infinite) T-graph corresponding to probabilities p_a, p_b, p_c (chosen as above) and call it T. Let ℓ be the linear map as described in Proposition 3.2, we recall that the map ψ from $\mathcal{H}^\dagger$ to T is almost equal to ℓ.

We now describe how to construct subdomains U^δ of the scaled lattice $\delta\mathcal{H}$ together with a marked point x^δ which approximate (U, x) in the sense that the boundary of U^δ approximates ∂U as closed sets and that they have a marked boundary face x^δ such that $x^\delta \to x$. They also satisfy the following properties. Firstly, when viewed as a lozenge tiling of a subdomain (and thus equivalently as a stack of cubes in 3-d), the cubes on the boundary will be within $O(1)$ of some fixed plane $\mathcal{P}$. Another way of putting it is that if p_a, p_b, p_c are such that (p_a, p_b, p_c) is normal to the plane $\mathcal{P}$ then the height function of any dimer configuration on U^δ, measured with respect to the flow ϕ_{p_a,p_b,p_c}, will be within $O(1)$ from $\mathcal{P}$ along the boundary. Actually using Theorem 4.4, the height differences with respect to x^δ along the boundary are also given by the winding of the boundary of U^δ. Secondly, the dimer configuration will correspond to a wired UST by a finite analogue of the construction introduced in Section 4.1.

In what follows $o_\delta(1)$ denotes a function which goes to 0 as $\delta \to 0$.

Write $\tilde{U}$ for $\ell(U)$ and fix a conformal map $\varphi : \mathbb{D} \mapsto \tilde{U}$. We recall that φ extends continuously to $\bar{\mathbb{D}}$ if and only if $\partial\tilde{U}$ is locally connected, by Theorem 2.1 of [10]. We will assume that U and thus $\tilde{U}$ has locally connected boundary throughout this section. Hence ∂U and $\partial\tilde{U}$ can each be seen as a curve[2]. Consider $\tilde{U}_\varepsilon := \varphi((1-\varepsilon)\mathbb{D})$ (note that $\tilde{U}_\varepsilon \subset \tilde{U}$). Using the above information, we deduce that $\partial\tilde{U}_\varepsilon \to \partial\tilde{U}$ in uniform norm between curves up to reparametrisation[3].

Now we claim that one can find a sequence of directed simple loops in the scaled T-graph δT such that they converge to $\partial\tilde{U}$ in the uniform norm up to reparametrisation. To see this, consider the image of $T^\delta := \delta T \cap \tilde{U}$ under φ^{-1}. We will first find a path (in fact, an oriented loop) on $\varphi^{-1}(T^\delta) \subset \mathbb{D}$ such that the path will remain at distances between $1-2\varepsilon$ and $1-\varepsilon/2$ from the origin. Indeed, we can approximate the circle $(1-\varepsilon)\partial\mathbb{D}$ by a sequence of rectangles of dimension $\varepsilon/10 \times 3\varepsilon/10$ lying strictly inside $(1-\varepsilon/2)\partial\mathbb{D}$ such that the starting ball of one rectangle is the target ball of another one. We can use the crossing assumption on T and the bound on the derivative of $|\varphi'|$ to say that these rectangles also satisfy the crossing assumption on the graph $\varphi^{-1}(T^\delta)$ for small enough $\delta(\varepsilon)$. Indeed, for each such rectangle R with starting ball B_1 and target ball B_2, one can find a sequence of rectangles $R_1, \dots, R_k \subset \varphi(R)$ in $\tilde{U}$ each with aspect ratio $3:1$ such that the starting ball of R_1 contains $\varphi(B_1)$ and the target ball of R_k contains $\varphi(B_2)$ and the starting ball of R_i is the target ball of R_{i-1} for $i = 1, \dots, k$. Observe that the size of R_i can be chosen to be at least $c(\varepsilon)$ where $c(\varepsilon) > 0$ depends only on ε and U (in fact, on the minimal value of $|\varphi'|$ in $(1-\varepsilon/4)\mathbb{D}$). Using the crossing estimate, for small enough $\delta = \delta(\varepsilon)$, there exists at least one path crossing each rectangle $R_i, 1 \le i \le k$. These paths can be concatenated. Consequently we obtain a loop in T^δ which is formed by concatenation of paths lying inside the images of the rectangles in $\mathbb{D}$. We can assume that this path is a directed simple loop by loop-erasing and note that by construction, this still keeps some portion of the loop in $\varphi(R)$ for each rectangle R in $\mathbb{D}$. Call this loop C^δ. Note that $\varphi^{-1}(C^\delta)$ is a simple loop in $\mathbb{D}$ which lies between $(1-\varepsilon/2)\partial\mathbb{D}$ and $(1-2\varepsilon)\partial\mathbb{D}$ and hence $C^\delta \to \partial\tilde{U}$ in the uniform norm when parametrised by the argument as $\varepsilon \to 0$ since φ extends continuously to $\partial\mathbb{D}$.

Our second claim is that we can find a simple path $\tilde{P}^\delta$ starting from a point on C^δ within $o_\delta(1)$ distance of $\ell(x)$, and such that $\tilde{P}^\delta$ goes off to infinity avoiding C^δ except for the starting point. To see this the crucial point again is that we have a bound of the derivative of φ on a neighbourhood of C^δ.

Let us give a detailed explanation. We move back to the unit disc and assume without loss of generality that $\varphi^{-1}(x)$ is 1. From the bound on the derivative, note that every point in $\varphi(\mathbb{D}) = \tilde{U}$ outside $\tilde{U}_{1-\varepsilon/100} = \varphi((1-\varepsilon/100)\mathbb{D})$ is at least c_ε away from C^δ. Furthermore, consider the rectangle $R = [1-10\varepsilon, 1-\varepsilon/200) \times [-\varepsilon/200, \varepsilon/200] \subset \mathbb{D}$. By planarity and the crossing estimate there exists a simple path $p \subset \varphi(R)$ in the graph T^δ, connecting a vertex in $C^\delta \cap \varphi(R)$ to a vertex outside $\varphi((1-\varepsilon/100)\mathbb{D})$.

[2] i.e. as a possibly non-injective continuous mapping $\gamma : \mathbb{S}_1 \mapsto \mathbb{C}$

[3] In a non-locally connected boundary, like the topologist's comb, this convergence might fail.

The curve C^δ is Jordan and therefore defines an inside I^δ and an outside. Call y^δ the end point of the path p constructed above. Note that since $y^\delta \in \mathbb{C} \setminus I^\delta$ we can find a continuous path γ which connects y^δ to infinity, avoiding $\bar{I}^\delta$. By compactness, note that γ is at distance at least c'_ε from $\bar{I}^\delta$. Thus using a suitable sequence of rectangles of size $c'_\varepsilon/2$ and using the crossing estimate, we can find an infinite oriented simple path p' in T^δ which starts from y^δ and avoids C^δ. We define $\tilde{P}^\delta$ to be the portion of the path $p \cup p'$ from the last time it crosses C^δ. The point where $p \cup p'$ crosses C^δ for the last time is within $o_\varepsilon(1)$ of the starting point of p by the argument in the previous paragraph and hence within $o_\delta(1)$ of x.

We root C^δ at the vertex x^δ. Recall that C^δ is a rooted oriented cycle. Let P^δ be the path obtained from C^δ deleting the first edge of C^δ and concatenating with $\tilde{P}^\delta$. Thus P^δ is a simple oriented path going to infinity (which roughly speaking starts by looping around $\tilde{U}$ and then follows $\tilde{P}^\delta$). Call the erased edge e^δ. Since each vertex of the T-graph corresponds to a face of the hexagonal lattice $\mathcal{H}^\dagger$, $\psi^{-1}(C^\delta)$ is a self avoiding cycle in $\delta\mathcal{H}^\dagger$ by Proposition 3.4. We define $\tilde{U}^\delta$ to be the subgraph of $\delta\mathcal{H}$ spanned by vertices strictly inside that cycle. The dual edge of $\psi^{-1}(e^\delta)$ is an edge of $\delta\mathcal{H}$ with exactly one of its endpoints in $\tilde{U}^\delta$ and we define U^δ by removing this vertex from $\tilde{U}^\delta$. Define the boundary faces of U^δ to be the faces in $\psi^{-1}(C^\delta)$. (Note that it may happen that x^δ, although a boundary face of U^δ in this definition, is not actually adjacent to U^δ after the vertex removal, but is always within $O(\delta)$ from U^δ, see Figure 10)

Now the idea will be to use the connection between dimers and trees in the whole plane and to that end we will extend P^δ to a one ended tree in the full plane.

Let D^δ be the sub-graph of δT consisting of all the vertices which are strictly inside C^δ and all the outgoing edges from them (including their endpoints). Clearly, every vertex in D^δ has outdegree 0 or 2. Again, using the crossing estimate, P^δ can be extended to a one-ended spanning tree of T^δ. Furthermore the extensions inside and outside C^δ do not depend on each other (in the sense that given an extension outside, every extension inside is possible) because the edge e^δ is never used. The possible extensions in D^δ are exactly the wired trees of D^δ (where we glue together all the degree 0 vertices).

Let F be an arbitrary extension of P^δ into a one-ended tree in δT. Notice that $F^\dagger$ can be decomposed into $F^\dagger = F^\dagger_{\text{int}} \sqcup F^\dagger_{\text{ext}} \sqcup (e^\delta)^\dagger$, where $F^\dagger_{\text{int}}$ is the portion of $F^\dagger$ spanned by vertices lying completely inside C^δ, $F^\dagger_{\text{ext}}$ is the portion of $F^\dagger$ spanned by vertices lying strictly outside C^δ and $(e^\delta)^\dagger$ is the dual edge to e^δ. Furthermore $F^\dagger_{\text{ext}}$ and $F^\dagger_{\text{int}}$ are joined together by $(e^\delta)^\dagger$ oriented from $F^\dagger_{\text{int}}$ to $F^\dagger_{\text{ext}}$. Hence the oriented tree $F^\dagger_{\text{int}}$ is a function of the oriented tree F_{int}.

Recall the map M from Definition 4.1 which produces a matching from a one-ended spanning tree. $M(F^\dagger)$ is a perfect matching of the whole hexagonal lattice (by Proposition 4.2). Let h be the associated height function as in Theorem 4.4, defined up to a global additive constant. Note that $F^\dagger$ contains no edge connecting the inside of C^δ with the outside of C^δ (except for $(e^\delta)^\dagger$). Therefore $M(F^\dagger_{\text{int}})$ is a perfect matching of U^δ. It then follows from Theorem 3.4 in [6] (applied to $F^\dagger_{\text{int}}$) that the image of the wired UST measure on D^δ is the uniform dimer measure on U^δ. Furthermore, for any $v \in \partial U^\delta$, $h(v) - h(x^\delta)$ is the intrinsic winding of the oriented path from $\psi(x^\delta)$ to $\psi(v)$ in C^δ (in the sense of Theorem 4.4). Let us summarise our findings in the following proposition.

Theorem 4.8. *The objects $U^\delta, D^\delta, x^\delta, e^\delta$ constructed above satisfy the following:*

- *The graph U^δ is a subgraph of δH. Furthermore, the boundary faces of U^δ (faces which have neighbour outside) can be joined together in the dual graph to form a simple loop and this loop approximates U in the uniform norm as unparametrised curved with $x^\delta \to x \in \partial U$. Further D^δ is a subgraph of T^δ.*
- *The images by the construction of Definition 4.1 of the wired UST in D^δ with the dual tree oriented towards e^δ, are uniform dimer configurations U^δ.*
- *For any $v \in \partial U^\delta$, $h(v) - h(x^\delta)$ is the intrinsic winding of the boundary of D^δ between $\psi(v)$ and $\psi(x^\delta)$ (more precisely, it is the winding of P^δ constructed above between $\psi(v)$ and $\psi(x^\delta)$ in the sense of Theorem 4.4).*

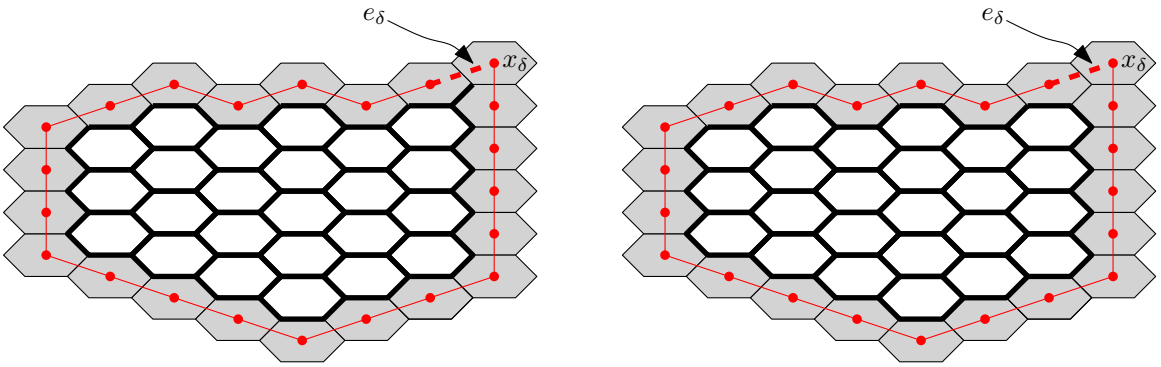

Figure 10: Illustration of the construction of the domain U^δ. The red loop (corresponding to the grey hexagons) is $\psi^{-1}(C^\delta)$. The dotted red edge is e^δ. On the left, the bold hexagons form $\tilde{U}^\delta$. On the right we delete a vertex from $\tilde{U}^\delta$ to obtain U^δ. In this case, x^δ is not actually adjacent to U^δ.

Remark 4.9. Given our choice of reference flow, the third point in Theorem 4.8 together with Remark 4.7 implies that the height function on U^δ, viewed as a surface in $2+1$ dimensions via cube stacks, lies within $O(1)$ of the plane $\mathcal{P}$. The domains of U^δ constructed here can be thought of as a natural generalisation of **Temperleyan domains** in the case of arbitrary slope and to the hexagonal lattice.

Note that there exist subsets of the hexagonal lattice that admit dimer configurations but are not Temperleyan in the above sense.

4.3 Local limits

In this section we show that the mapping M between a full plane uniform spanning tree measure in a *translation invariant version* of the T-graph corresponding to a certain slope produces a translation invariant ergodic Gibbs measure on dimer configurations with the same slope. The translation invariant version amounts to randomising the T-graph by picking a λ which is uniform on the unit circle as explained in Proposition 3.1. To emphasise a bit more, it is straightforward to see from locality of the map M that once we establish a local limit for the trees the corresponding dimer configuration is Gibbs. Also since the T-graph is translation invariant, so is the dimer measure. However, translation invariant Gibbs measures are not unique and we need an extra argument to identify the limit.

Note that it is not obvious from classical electrical network argument that the local limit of wired UST exists since our graph is oriented. However we have established this in Corollary 4.20 in [1], using the uniform crossing estimate of Theorem 3.8.

Theorem 4.10. *Fix $p_a, p_b, p_c > 0$ be such that $p_a + p_b + p_c = 1$. Let λ be picked uniformly at random from the unit circle. Let $T = T_{\lambda\Delta}$ be a (random) T-graph associated with λ and p_a, p_b, p_c as in Proposition 3.1. Let μ_F denote the uniform spanning tree measure on T defined by taking a local limit of wired domains as in Corollary 4.20 in [1]. Then the random dimer configuration $M(\mu_F)$ (as in Definition 4.1) is the unique translation invariant ergodic Gibbs measure on full plane lozenge tiling with probabilities p_a, p_b, p_c.*

Proof. First, condition on λ and take a sequence of domains as described in Theorem 4.8 which exhaust the plane. Thus there is a measure preserving transformation between uniform spanning trees in this sequence and the uniform dimer configurations in these domains (i.e. with weight 1 for each dimer). Since the transformation going from tree to a dimer configuration is local, we take limits of the uniform spanning trees via Corollary 4.2 in [1] and obtain that $M(\mu_F)$ is a Gibbs measure on lozenge tilings of the full plane. We now average over λ to obtain a translation invariant Gibbs measure on dimer configurations which we write

μ_{dim}. We can decompose μ_{dim} into its ergodic parts which are fully characterised by Theorem 2.1:

$$\mu_{\text{dim}} = \int \mu_s d\nu(s)$$

for some measure ν on $S = \{s = (p_a, p_b, p_c) \in [0,\infty)^3 : p_a + p_b + p_c = 1\}$.

Let $s_0 = (p_a, p_b, p_c)$ denote the triplet used to construct the T-graph. Recall that if we measure height functions using the reference flow defined from the T-graph (see Remark 4.7), we have $\mathbb{E}_{\mu_{s_0}}[h(x)] = O(1)$ for all $x \in \mathcal{H}^\dagger$ and $h(0) = 0$ (in the full plane we always need to pin the height at one point). For other measures it is also clear by Theorem 2.1 that

$$\mathbb{E}_{\mu_s}[h(x)] - \mathbb{E}_{\mu_{s_0}}[h(x)] = l_s(x)$$

where l_s is a non-zero linear function. Furthermore it is known that for any $s \in S$, the fluctuations around the mean of the height is at most $L^{1/4}$ with very high probability (in fact the right order is logarithmic, but not essential for this proof). For instance, it follows from [9], Theorem 2.8 for all $L > 0$, for all $x \in B(0, L) \cap \mathcal{H}^\dagger$,

$$\mathbb{P}_{\mu_s}(|h(x) - \mathbb{E}_{\mu_s}(h(x))| > L^{1/4}) < L^{-10}. \tag{4.1}$$

Let us assume by contradiction the measure ν is not a Dirac mass at s_0. Consider the domain U approximating $B(0, L)$ as in the construction of Theorem 4.8 for a large radius L. On one hand, Theorem 2.9 in [9] shows that under the uniform measure on U,

$$\mathbb{P}\Big(\sup_{x \in B(0,L)} |h(x)| > L^{1/4}\Big) \leq L^{-10}, \tag{4.2}$$

where the height $h(x)$ is measured with respect to the reference flow on the T-graph. Since ν is not a Dirac mass at s_0, we can find some $\varepsilon > 0$ such that the ν measure of

$$\Sigma := \{s : \exists x, l_s(x) \geq 2\varepsilon |x|\}$$

is at least 4ε. By (4.1), for any L large enough we have

$$\forall s \in \Sigma, \quad \mathbb{P}_{\mu_s}\left(\sup_{x \in B(0,\sqrt{L})} |h(x)| \geq \varepsilon\sqrt{L}\right) \geq 1 - \varepsilon,$$

where h is still measured with respect to the reference flow. Along with the bound on $\nu(\Sigma)$, this implies

$$\mathbb{P}_{\mu_{\text{dim}}}\left(\sup_{x \in B(0,\sqrt{L})} |h(x)| \geq \varepsilon\sqrt{L}\right) \geq 3\varepsilon.$$

Now by the coupling of Corollary 4.2 in [1], we can couple the dimer configuration in $B(0, \sqrt{L})$ under the uniform measure μ_U in U and the whole plane measure μ_{dim} with probability at least $1 - \varepsilon$ (actually this was shown even conditioned on λ). In particular

$$\mathbb{P}_{\mu_U}(\sup_{x \in B(0,\sqrt{L})} |h(x)| \geq \varepsilon\sqrt{L}) \geq \mathbb{P}_{\mu_{\text{dim}}}\Big(\sup_{x \in B(0,\sqrt{L})} |h(x)| \geq \varepsilon\sqrt{L}\Big) - \varepsilon \geq 2\varepsilon.$$

This is a contradiction with (4.2). □

References

[1] N. Berestycki, B. Laslier, and G. Ray. Universality of fluctutations in the dimer model. *arXiv:1603.09740*, 2016.

[2] R. Kenyon. Lectures on dimers. IAS/Park City mathematical series, vol. 16: Statistical mechanics, ams, 2009. *arXiv preprint arXiv:0910.3129*.

[3] R. Kenyon. Height fluctuations in the honeycomb dimer model. *Communications in Mathematical Physics*, 281(3):675–709, 2008.

[4] R. Kenyon, A. Okounkov, and S. Sheffield. Dimers and amoebae. *Annals of mathematics*, pages 1019–1056, 2006.

[5] R. W. Kenyon, J. G. Propp, and D. B. Wilson. Trees and matchings. *Electron. J. Combin.*, 7:Research Paper 25, 34 pp. (electronic), 2000.

[6] R. W. Kenyon and S. Sheffield. Dimers, tilings and trees. *J. Combin. Theory Ser. B*, 92(2):295–317, 2004.

[7] B. Laslier. Central limit theorem for T-graphs. *arXiv preprint arXiv:1312.3177*, 2013.

[8] B. Laslier and F. L. Toninelli. Lozenge tilings, glauber dynamics and macroscopic shape. *Comm. Math. Phys.*, 2013. To appear, arXiv:1310.5844.

[9] B. Laslier and F. L. Toninelli. How quickly can we sample a uniform domino tiling of the $2L \times 2L$ square via Glauber dynamics? *Probab. Theory Related Fields*, 161(3-4):509–559, 2015.

[10] C. Pommerenke. *Boundary behaviour of conformal maps*, volume 299 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, 2013.

[11] S. Sheffield. Random surfaces. *Astérisque*, (304):vi+175, 2005.